

Measuring Skin Reflectance and Subsurface Scattering

Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Addy Ngan, Markus Gross

TR2005-046 December 2005

Abstract

It is well known that human facial skin has complex reflectance properties that are difficult to model, render, and edit. We built two measurement devices to analyze and reconstruct the reflectance properties of facial skin. One device is a dome-structured face scanning system equipped with 16 cameras and 150 point light sources that is used to acquire BRDF samples of a face. The other device is a touch-based HDR imaging system to measure subsurface scattering properties of the face. In this technical report we describe how these devices are constructed, calibrated, and used to acquire high-quality reflectance data of human faces.

SIGGRAPH 2005

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Measuring Skin Reflectance and Subsurface Scattering

Tim Weyrich* Wojciech Matusik† Hanspeter Pfister‡ Addy Ngan‡ Markus Gross*

Abstract

It is well known that human facial skin has complex reflectance properties that are difficult to model, render, and edit. We built two measurement devices to analyze and reconstruct the reflectance properties of facial skin. One device is a dome-structured face scanning system equipped with 16 cameras and 150 point light sources that is used to acquire BRDF samples of a face. The other device is a touch-based HDR imaging system to measure subsurface scattering properties of the face. In this technical report we describe how these devices are constructed, calibrated, and used to acquire high-quality reflectance data of human faces.

1 Face Scanning Dome

Figure 1 shows a photograph of our face scanning dome. The sub-




Figure 1: The face scanning dome consists of 16 digital cameras, 150 LED light sources, and a commercial 3D face scanning system.

ject sits in a chair with a head rest to keep the head still during the capture process. The chair is surrounded by 16 cameras and 150 LED light sources that are mounted on a geodesic dome. The dome is a once-subdivided icosahedron with 3 meters diameter, 42 vertices, 120 edges, and 80 faces. Some edges in the back of the dome were left out to leave an opening for a person. The one bottom ring was also left out so that the head rest of the chair is approximately in the center of the dome. Placing two light sources on each remaining edge allows 150 lights to be evenly distributed on the sphere. Our dome is closely related to the Light Stage 3 system by Debevec et al. [Debevec et al. 2002].

Each light is custom built and contains 103 white LEDs. They are mounted in a 15 cm diameter cylinder and covered by a glass diffuser (see Figure 2). We experimented with the iColor MR lights from Color Kinetics corporation. However, their mixture of red, green, and blue LEDs led to noticeable spectral artifacts. The spectrum of white LEDs is not even, and the light power shows consid-




Figure 2: Two of the 150 custom built lights mounted in the dome. The diffuser on the right was removed to show the arrangement of the 103 white LEDs.

erable falloff across the beam. We discuss our correction for both effects in Section 1.1. The lights are connected to a common power source. The light intensity is controlled by modulating the pulse width of the current to the LEDs.

Each vertex in the dome holds a Bogen adjustable camera mount. We are using 16 Basler A101fc color cameras with 1300×1030 8-bit per pixel CCD sensors. The positions of the cameras are chosen to cover the hemisphere in front of the subject as regularly as possible. Two cameras each are connected to one of eight PCs via FireWire bus. The maximum frame rate of the cameras is 12 frames per second. The system records the raw images to 2 GB of RAM in each PC before storing them to disk. During post-processing we use factory supplied software for color interpolation from the raw Bayer pattern.

The cameras and lights are connected with individual control wires to a custom-built PCI board inside a PC. A series of custom programmable logic devices (CPLD) precisely trigger the cameras and lights under software control. In this application, the system sequentially turns each light on while simultaneously capturing images with all 16 cameras. To deal with the large intensity differences between specular highlights and shadows we capture high-dynamic range (HDR) images [Debevec and Malik 1997]. We immediately repeat the capture sequence with two different exposure settings using an increment of 1.4 f-stops.

The complete sequence takes about 25 seconds for the two passes through all 150 light sources (limited by the frame rate of the cameras) and collects 4800 reflectance images of a subject's face. Our system is also set to collect two additional images of the subject per each camera and each exposure setting just before and after the lighting sequence. These *black* images are used to subtract the ambient illumination as well as the DC offset of the image sensor and the AD converter. Figure 3 shows some sample reflectance images taken from all 16 cameras in the higher exposure setting.

To detect head movement we perform an analysis of the images right after capture. Pairs of images were acquired about 12 seconds apart with the same camera / light combination and different exposure settings. We estimate the head translation by computing the autocorrelation between five of these image pairs. If any of them show too much translation we are repeating the scan. The image pairs are chosen from orthogonal views, so even rotational head motion will produce a virtual translation in at least one of the views. The remaining head motion or slight changes in facial expressions lead to noise in our reflectance measurements.

A commercial face scanning system from 3QTech is placed behind openings of the dome. Using two structured light projectors and four cameras it captures the complete 3D face geometry in less

*ETH Zürich, Email: [weyrich,grossm]@inf.ethz.ch

†MERL, Email: [matusik,pfister]@merl.com

‡MIT, Email: addy@graphics.csail.mit.edu

than one second. The output mesh contains about 40,000 triangles and resolves features as small as 1 mm.

1.1 Dome Calibration

We use the freely available OpenCV calibration package to determine spherical and tangential lens distortion parameters. External camera calibration is performed by synchronously acquiring a few hundred images of an LED swept through the dome center. Non-linear optimization is used to optimize for the remaining camera intrinsics and extrinsics.

The geometry from the 3QTech scanner needs to be registered with the image data. We scan a calibration target with the 3QTech scanner and take images with all cameras under ambient illumination. Point correspondences between the images are used to reconstruct the 3D positions of nine feature points on the target. The Procrustes method is used to find a rigid transformation and scale to match the corresponding points on the scanned 3D geometry. We then compute the transformation between scanner and camera coordinates. We also register the camera and light source positions with a CAD model of the dome to get the transformations to the global dome coordinate system.

Camera vignetting leads to an intensity fall-off towards image borders. We calibrate for it by acquiring images of a sheet of vellum, back-lit by ambient illumination. To remove small-scale variations of the calibration vellum, we fit a bi-variate low-degree polynomial to the observed intensity fall-off. We equalize the spectral characteristics of our cameras by acquiring images of a color chart under controlled lighting conditions. We found that an affine transformation of color values is sufficient to match the different response curves of our cameras.

We also perform calibration of the light flux and beam spread. For each light source, we take images of a 12" x 12" Fluorilon reflectance target with calibration markers under different orientations. The target diffusely reflects 99.9% of the incident light, yielding hundreds of thousands reflectance samples per light source. The light source diffusers were chosen to produce a very smooth beam fall-off, and the cross-section of the beam spread can be well approximated by a bi-variate 2nd degree polynomial. Because of the large distance between light sources and the face we ignore near-field effects.

2 Subsurface Measurement Device

Our subsurface measurement device is an image-based version of a fiber optic spectrometer with a linear array of optical fiber detectors [Nickell et al. 2000] (see Figures 4 and 5). White light from




Figure 4: A schematic of our subsurface measurement device.



Figure 5: Left: A picture of the sensor head with linear fiber array. The source fiber is lit. Right: The fiber array leads to a camera in a light-proof box. The box is cooled to minimize imaging sensor noise.

an LED (the same kind that is used for the dome lights) is coupled to a source fiber. The alignment of the fibers is linear to minimize sensor size. A sensor head holds the source fiber and 28 detection fibers. Each fiber has a core diameter of 980 μm and a numerical aperture of 0.51. A digital camera records the light collected by the detector fibers. We use a QImaging QICAM camera with 1360 x 1036 10-bit per pixel CCD sensor. The camera and detector fibers are encased in a light-proof box with air cooling to minimize imager noise. We capture 23 images bracketed by 2/3 f-stops to compute an HDR image of the detector fibers. The total acquisition time is about 88 seconds.

Figure 6 shows the sensor head placed on a face. We found that




Figure 6: Left: The sensor head placed on a face. Top: Sensor fiber layout. The source fiber is denoted by a cross. Bottom: An HDR image of the detector fibers displayed with three different exposure values.

pressure variations on the skin caused by the mechanical movement of the sensor head influence the results. To maintain constant pressure between skin and sensor head we attached a silicon membrane connected to a suction pump. This greatly improves the repeatability of the measurements.

The detector distances from the source fiber are between 0.33 mm and 25 mm, and the spacing between fibers is 0.33 mm (see Figure 6). The placement of detector fibers is asymmetric (more fibers are placed on one side of the source fiber) and slightly staggered (for tighter spacing). Small spacers (indicated by little rectangles in the figure) offset the fibers on one side by 0.5 mm. This leads to denser measurements near the point of light injection. To obtain two-dimensional data we rotate the sensor head around the source fiber in steps of 10 degrees. This was achieved by marking two end points of the sensor head on the skin with a dry-erase marker.

Due to the relatively large size of the sensor head we are not able to measure regions with large curvature (e.g., nose or chin). And for hygienic reasons we do not measure lips (although, as expected, some experiments showed large variations between skin and lips). We found that diffuse subsurface scattering varies smoothly within the face, allowing us to interpolate data from only a few measurement points. We have chosen to measure subsurface scattering on forehead, cheek, and below the chin.




Figure 3: Raw reflectance images of a subject acquired with all 16 cameras and 14 (of 150) lighting conditions. Each row shows images with different camera viewpoints and the same lighting condition. Each column displays images with different lighting conditions and the same viewpoint. The displayed images were all captured using the higher exposure setting.

2.1 Subsurface Device Calibration

Before we can proceed with measurements we need to compensate for the different transmission characteristics between detection fibers. We measure uniform irradiance from a lightbox with the probe head. Based on these measurements we compute appropriate normalization factors for the different detection fibers. To detect potential changes in the intensity of the LED light source, a calibration fiber leads directly from the light source to the camera (see Figure 4).

Similar to the acquisition method of Jensen et al. [2001], we model the influence of the angle of observation and Fresnel terms of our detection fibers by a constant factor K . To determine K we use skim milk as our calibration standard. Its scattering properties have been measured by Jensen et al. [2001], and – unlike other calibration materials – we achieved reproducible results with skim milk.

References

- DEBEVEC, P., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Computer Graphics, SIGGRAPH 97 Proceedings*, 369–378.
- DEBEVEC, P., WENGER, A., TCHOU, C., GARDNER, A., WAESE, J., AND HAWKINS, T. 2002. A lighting reproduction approach to live-action compositing. *ACM Transactions on Graphics (SIGGRAPH 2002)* 21, 3 (July), 547–556.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Computer Graphics, SIGGRAPH 2001 Proceedings*, 511–518.
- NICKELL, S., HERMANN, M., ESSENPREIS, M., FARRELL, T. J., KRAMER, U., AND PATTERSON, M. S. 2000. Anisotropy of light propagation in human skin. *Phys. Med. Biol.* 45, 2873–2886.